A Proper Method for Introducing Shear into Compressible RANS Models

M. Ulitsky

July 15, 2014

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

A Proper Method for Introducing Shear into Compressible RANS Models

Mark Ulitsky and Brian Yang
AX Division
Lawrence Livermore National Security, LLC.

Contents

1 Introduction

A very generic way of writing a transport equation for k (the turbulent kinetic energy per unit mass) for a variable density flow is of the form:

$$\frac{D\overline{\rho}k}{Dt} = S_{RM} + S_{RT} + S_{KH} + S_{dilatation} + S_{diffusion} - S_{dissiaption}, \tag{1}$$

where the subscripts on the source terms refer respectively to the Richtmeyer-Meshkov, Rayleigh-Taylor, and Kelvin-Helmholtz instabilities. That is, we tend to think of splitting the turbulence into the various hydrodynamic assaults that it is likely to encounter in a compressible flow. Unfortunately, there is a subtle issue related to double counting among the shear and dilatation terms that does not occur for incompressible flows, but can have serious consequences for their compressible counterparts.

Dilatation (expansion/compression) in RANS models was introduced more as an afterthought, rather than as an intentional physical phenomenon that required modeling. For example, it is very common to represent shear or the Kelvin-Helmholtz instability by using a Boussinesq approximation for R_{ij} (the turbulent Reynolds stress) that contains a turbulent pressure part (P_{ij}) and a deviatoric part (D_{ij}) of the form:

$$R_{ij} \equiv \overline{\rho u_i'' u_j''} = P_{ij} - D_{ij} \tag{2}$$

$$P_{ij} = \frac{2}{3}\overline{\rho}k\delta_{ij} \tag{3}$$

$$D_{ij} = 2\mu_t\left(S_{ij} - \frac{1}{3}\nabla \cdot \tilde{\mathbf{u}}\,\delta_{ij}\right) \tag{4}$$

$$S_{ij} = \frac{1}{2}\left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i}\right), \tag{5}$$

where δ_{ij} is the Kronecker delta tensor, S_{ij} is the rate of strain tensor (the symmetric part of the velocity gradient tensor), and μ_t, the turbulent viscosity is given by $C_\mu\overline{\rho}k^2/\epsilon$ for a k-ϵ model and by $C_\mu\overline{\rho}\sqrt{k}L$ for a k-L model. Also, in the above equations, we are using the standard nomenclature convention of an overbar to represent an ensemble average, a tilde to represent a Favre average, and double primes to represent a fluctuation from the Favre mean.

The transport equation for the momentum will have a term like $\nabla \cdot \mathbf{R}$ on the right hand side and the k equation will have a term like $\mathbf{R} : \mathbf{S}$, which expands to $P_{ij}S_{ij}$ and $-D_{ij}S_{ij}$. Using the above definitions, we see that the dilatation term in the k equation is given by

$$P_{ij}S_{ij} = \frac{2}{3}\overline{\rho}k\,\nabla \cdot \tilde{\mathbf{u}}, \tag{6}$$

and the shear term is given by $-D_{ij}S_{ij}$.

Since deviatoric tensors are traceless by definition, it was required for compressible flow scenarios to express D_{ij} as

$$D_{ij} = 2\mu_t\left(S_{ij} - \frac{1}{3}\nabla \cdot \tilde{\mathbf{u}}\,\delta_{ij}\right), \tag{7}$$

where the second term is exactly what is required to ensure that $D_{ii} = 0$. There are several well-known problems with the Boussinesq form for D_{ij} that are related to tensor unrealizability. That is, if D_{ij} gets too large as compared to P_{ij}, then the diagonal elements of R_{ij} can become negative and the off-diagonal elements can violate the Cauchy-Schwartz condition. The determinant of R_{ij} can also become negative. These instances of unrealizability or violations of the positive definiteness of the stress tensor predominantly occur in shocks and strong rarefactions, and if realizability is not enforced (typically this is done by scaling down the deviatoric tensor), then spurious production of k will result.

What has not been discussed previously in the literature is that the Boussinesq shear approximation for D_{ij} is NOT free of compression/expansion and so ALL RANS models that use this form are double counting dilatation effects. This double counting greatly exacerbates the tensor realizability problem. It also leads to many turbulence codes having switches like "turn off the deviatoric stress in a shock" or something similar. These switches are problematic for multiple reasons, the first being that there is no perfect method of doing shock detection in general, especially on unstructured meshes, where a single threshold for detecting a shock may not be appropriate as the mesh is refined. The second reason is that these types of switches or robustness fixes will not be necessary if one uses a form for D_{ij} that is specifically constructed to be well-behaved for compressible flows. Even worse, when a flow is truly in compression or expansion (and there is no shear), D_{ij} does not go to zero, as it should. The bottom line is that special care and thought needs to go into a model for D_{ij} in order to achieve reasonable behavior for compressible flow. No amount of clever tensor invariant manipulations will fix the problem either. In fact, just like the solution to the Reimann problem which requires several branches to distinguish between shocks, contacts. and rarefactions for the different left and right states, the solution to the double counting problem will involve branches that depend on the sign of $\nabla \cdot \tilde{\mathbf{u}}$ and the signs of the individual eigenvalues of S_{ij}. Thus, the solution cannot be expressed in simple tensor covariant form!

The rest of this paper will delve into the details about how to accomplish the important objective of eliminating double counting between shear and dilatation. Any fixup of the Boussinesq approximation for compressible flow that does not addresss the double counting issue is simply incorrect. That is, the double counting issue is of fundamental and paramount importance, and other corrections will only make sense in a compressible flow if the double counting problem is dealt with first.

2 Eigenvalues and eigenvectors

We have made the claim that the standard form for the Boussinesq model of D_{ij} is not well formulated for compressible flow. The goal of this section is to replace the equation given by

$$D_{ij} = 2\mu_t \left(S_{ij} - \frac{1}{3}\nabla \cdot \tilde{\mathbf{u}} \, \delta_{ij} \right) \tag{8}$$

with one of the form

$$D_{ij} = 2\mu_t S_{ij}^* , \tag{9}$$

where S_{ij}^* is compression/expansion free and only is active if there is true shear present in the flow.

It should not be too surprising that in order to properly decompose S_{ij} into a compression/expansion part and a true shear part, that we will need to make a transformation into the principal axes. Once we are in the principal coordinate system, we will apply a few reasonable constraints to allow us to uniquely determine the compression/expansion part. A simple subtraction will then result in the true shear part and the final step will be to take the shear part and transform it back into the original coordinate system.

The main advantage of working in the prinipal axes is that the tensor is diagonal in its proper frame. For example, if $\mathbf{A}$ is a symmetric second rank tensor, then the relation between $\mathbf{A}$ in the original and proper coordinate systems is given by:

$$
\begin{align}
\mathbf{A} &= A_{11}\boldsymbol{\delta}_1\boldsymbol{\delta}_1 + A_{12}\left(\boldsymbol{\delta}_1\boldsymbol{\delta}_2 + \boldsymbol{\delta}_2\boldsymbol{\delta}_1\right) \tag{10}\\
&+ A_{22}\boldsymbol{\delta}_2\boldsymbol{\delta}_2 + A_{23}\left(\boldsymbol{\delta}_2\boldsymbol{\delta}_3 + \boldsymbol{\delta}_3\boldsymbol{\delta}_2\right) \\
&+ A_{33}\boldsymbol{\delta}_3\boldsymbol{\delta}_3 + A_{13}\left(\boldsymbol{\delta}_1\boldsymbol{\delta}_3 + \boldsymbol{\delta}_3\boldsymbol{\delta}_1\right) \\
&= \lambda_1\mathbf{e}_1\mathbf{e}_1 + \lambda_2\mathbf{e}_2\mathbf{e}_2 + \lambda_3\mathbf{e}_3\mathbf{e}_3 , \tag{11}
\end{align}
$$

where the subscripts '1', '2', and '3' in the original coordinate system could stand for 'x', 'y', and 'z', or 'r', 'θ', 'z', and λ_i and $\mathbf{e}_i$ are the real eigenvalues and unit eigenvectors of S_{ij}. The advantages of working in the proper coordinate system should now be readily apparent.

The definition of shear in the principal axes is very simple. It is given by a compression along one principal axis and an expansion of the same magnitude along a different axis. In effect, this reduces to pairs of 1 and -1 on the diagonal (if we non-dimensionalize). This means that to have shear in the first place, we must have both positive and negative eigenvalues. If all of the eigenvalues are positive, then we have a true expansion and no shear. If all of the eigenvalues are negative, then we have a true compression and no shear. If we have both positive and negative eigenvalues, then we will seek a decomposition of the eigenvalues into a dilatation and shear part of the form:

$$
\begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix} = \begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} + \begin{pmatrix} \text{shear}_1 & & \\ & \text{shear}_2 & \\ & & \text{shear}_3 \end{pmatrix} . \tag{12}
$$

Now we seek to place some constraints that will allow us to uniquely determine the dilatation part of the decomposition. Our goal is to adjust the λ_i in such a way that the sum of the changes is identically zero, and the changes themselves are "small".

$$
\begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} = \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix} + \begin{pmatrix} \Delta\lambda_1 & & \\ & \Delta\lambda_2 & \\ & & \Delta\lambda_3 \end{pmatrix} \tag{13}
$$

$$
\sum_i \Delta\lambda_i = 0 \tag{14}
$$

$$
|\Delta\lambda_i| \leq |\lambda_i| . \tag{15}
$$

The first constraint is necessary because $S_{ii} = \lambda_1 + \lambda_2 + \lambda_3 = \nabla \cdot \tilde{\mathbf{u}}$, and we don't want to change the value or sign of $\nabla \cdot \tilde{\mathbf{u}}$. The second constraint results from wanting to limit the changes so that the λ_i do not flip sign.

When there are both positive and negative eigenvalues, there are 4 cases that need to be considered. In what follows, we will assume that when $\nabla \cdot \tilde{\mathbf{u}} > 0$ we will sort the λ_i in decreasing order so that $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and when $\nabla \cdot \tilde{\mathbf{u}} < 0$, we will sort the λ_i in increasing order so that $\lambda_1 \leq \lambda_2 \leq \lambda_3$. This sorting procedure will allow us to re-use the same formulas when the dilatation changes sign. The different cases are given by

1. $\nabla \cdot \tilde{\mathbf{u}} > 0$ and $\lambda_2 < 0$

2. $\nabla \cdot \tilde{\mathbf{u}} > 0$ and $\lambda_2 > 0$

3. $\nabla \cdot \tilde{\mathbf{u}} < 0$ and $\lambda_2 < 0$

4. $\nabla \cdot \tilde{\mathbf{u}} < 0$ and $\lambda_2 > 0$

Let's consider the first case, where the sum of the eigenvalues is positive, $\lambda_1 > 0$, and λ_2, λ_3 are negative. See figure Fig. 1. For this particular case, we will drive the 2 negative eigenvalues to zero and move λ_1 down so the sum of the changes is zero. This leads to

$$
\Delta\lambda_1 = \lambda_2 + \lambda_3 , \quad \Delta\lambda_2 = -\lambda_2 , \quad \Delta\lambda_3 = -\lambda_3 . \tag{16}
$$

We can now express the dilatation part exactly and subtract to find the shear part.

$$\begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} = (\lambda_1 + \lambda_2 + \lambda_3) \begin{pmatrix} 1 & & \\ & 0 & \\ & & 0 \end{pmatrix} \tag{17}$$

$$\begin{pmatrix} \text{shear}_1 & & \\ & \text{shear}_2 & \\ & & \text{shear}_3 \end{pmatrix} = \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix} - \begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} \tag{18}$$

$$= \lambda_2 \begin{pmatrix} -1 & & \\ & 1 & \\ & & 0 \end{pmatrix} + \lambda_3 \begin{pmatrix} -1 & & \\ & 0 & \\ & & 1 \end{pmatrix}. \tag{19}$$

For the second case we still have that the sum of the eigenvalues is positive, but now, λ_1 and $\lambda_2 > 0$, while λ_3 is negative. See figure Fig. 5. For this particular case, we will drive the negative eigenvalue to zero and partition the downward movement of λ_1 and λ_2 by fractionally weighting λ_3. The idea being that if $\lambda_1 = \lambda_2$, we would not bias either eigenvalue, but rather would move both by the same amount. This leads to

$$\Delta\lambda_1 = \left(\frac{\lambda_1}{\lambda_1 + \lambda_2}\right) \lambda_3, \quad \Delta\lambda_2 = \left(\frac{\lambda_2}{\lambda_1 + \lambda_2}\right) \lambda_3, \quad \Delta\lambda_3 = -\lambda_3 . \tag{20}$$

Repeating the same procedure as before now gives

$$\begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} = \left(\frac{\lambda_1 + \lambda_2 + \lambda_3}{\lambda_1 + \lambda_2}\right) \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & 0 \end{pmatrix} \tag{21}$$

$$\begin{pmatrix} \text{shear}_1 & & \\ & \text{shear}_2 & \\ & & \text{shear}_3 \end{pmatrix} = \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix} - \begin{pmatrix} \text{dil}_1 & & \\ & \text{dil}_2 & \\ & & \text{dil}_3 \end{pmatrix} \tag{22}$$

$$= \left(\frac{\lambda_1 \lambda_3}{\lambda_1 + \lambda_2}\right) \begin{pmatrix} -1 & & \\ & 0 & \\ & & 1 \end{pmatrix} + \left(\frac{\lambda_2 \lambda_3}{\lambda_1 + \lambda_2}\right) \begin{pmatrix} 0 & & \\ & -1 & \\ & & 1 \end{pmatrix}.$$

There should be no concern in these equations of the denominator $(\lambda_1 + \lambda_2)$ going to zero, since by construction, λ_1 and λ_2 both have the same sign.

Of course, the goal is to have an expression for the shear part that we can apply in the original coordinate system. This is trivial however, since we are in a diagonal representation. For example, if a tensor $\mathbf{A}$ in the eigenbasis has the form

$$\begin{pmatrix} \alpha & & \\ & \beta & \\ & & \gamma \end{pmatrix} \tag{23}$$

then this is equivalent to

$$\mathbf{A} = \alpha \mathbf{e}_1 \mathbf{e}_1 + \beta \mathbf{e}_2 \mathbf{e}_2 + \gamma \mathbf{e}_3 \mathbf{e}_3 \tag{24}$$

in the original coordinate system. That is, one simply has to take the unit eigenvectors and perform the necessary outerproducts.

Thus, for the case where $\nabla \cdot \tilde{\mathbf{u}} > 0$, the correct form for the true shear part, S^*_{ij}, is given by

$$\mathbf{S}^* = 0 \qquad\qquad \text{if all } \lambda_i \text{ have the same sign} \tag{25}$$

$$\mathbf{S}^* = -(\lambda_2 + \lambda_3) \mathbf{e}_1 \mathbf{e}_1 + \lambda_2 \mathbf{e}_2 \mathbf{e}_2 + \lambda_3 \mathbf{e}_3 \mathbf{e}_3 \qquad\qquad \text{if } \lambda_2 < 0 \tag{26}$$

$$\mathbf{S}^* = -\left(\frac{\lambda_1 \lambda_3}{\lambda_1 + \lambda_2}\right) \mathbf{e}_1 \mathbf{e}_1 - \left(\frac{\lambda_2 \lambda_3}{\lambda_1 + \lambda_2}\right) \mathbf{e}_2 \mathbf{e}_2 + \lambda_3 \mathbf{e}_3 \mathbf{e}_3 \qquad \text{if } \lambda_2 > 0 , \tag{27}$$

6

where the λ_i are sorted in decreasing order.

By referring to Figs. 5 and 4, one can go through similar arguments and algebra to arrive at the following expression for S_{ij}^* when $\nabla \cdot \tilde{\mathbf{u}} < 0$.

$$\mathbf{S}^* = 0 \qquad\qquad\qquad\qquad\qquad \text{if all } \lambda_i \text{ have the same sign} \quad (28)$$

$$\mathbf{S}^* = -\left(\lambda_2 + \lambda_3\right)\mathbf{e}_1\mathbf{e}_1 + \lambda_2\mathbf{e}_2\mathbf{e}_2 + \lambda_3\mathbf{e}_3\mathbf{e}_3 \qquad\qquad \text{if } \lambda_2 > 0 \qquad\qquad\quad (29)$$

$$\mathbf{S}^* = -\left(\frac{\lambda_1\lambda_3}{\lambda_1 + \lambda_2}\right)\mathbf{e}_1\mathbf{e}_1 - \left(\frac{\lambda_2\lambda_3}{\lambda_1 + \lambda_2}\right)\mathbf{e}_2\mathbf{e}_2 + \lambda_3\mathbf{e}_3\mathbf{e}_3 \qquad \text{if } \lambda_2 < 0 \ , \qquad\quad (30)$$

where now the λ_i must be sorted in increasing order. Thus by sorting the eigenvalues in the appropriate order, we arrive at essentially the same formulas.

Some points to note about the above formulas for S_{ij}^* are the following. When $\nabla \cdot \tilde{\mathbf{u}} \to 0$, we have $\mathbf{S}^* \to \mathbf{S}$, where $\mathbf{S}$ is given by

$$\mathbf{S}^* = \lambda_1\mathbf{e}_1\mathbf{e}_1 + \lambda_2\mathbf{e}_2\mathbf{e}_2 + \lambda_3\mathbf{e}_3\mathbf{e}_3 \ . \tag{31}$$

This can be verified by obverving that when $\nabla \cdot \tilde{\mathbf{u}} = 0$, we have the relation that $\lambda_1 + \lambda_2 + \lambda_3 = 0$. Of course this also follows from the fact that the dilatation part will vanish as $\nabla \cdot \tilde{\mathbf{u}} \to 0$, The above formulas for S_{ij}^* are also continuous as $\lambda_2 \to 0$. This is important, as it is the sign of λ_2 that determines whether we have 2 positive or 2 negative eigenvalues (when eigenvalues are present with different signs), and the formulas branch at this point.

3 Problems with Boussinesq

Now that we understand a proper way to express true shear in compressible flows, we are in a position to look more closely at the Boussinesq approximation. Essentially, the Boussinesq approximation models shear as:

$$\left(S_{ij} - \frac{1}{3}\nabla \cdot \tilde{\mathbf{u}}\,\delta_{ij}\right) \ . \tag{32}$$

This form can also be expressed as

$$\mathbf{S}^* = \left(\lambda_1 - \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}\right)\mathbf{e}_1\mathbf{e}_1 + \left(\lambda_2 - \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}\right)\mathbf{e}_2\mathbf{e}_2 + \left(\lambda_3 - \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}\right)\mathbf{e}_3\mathbf{e}_3 \ . \tag{33}$$

Since the Boussinesq approximation only knows about $\nabla \cdot \tilde{\mathbf{u}}$ and not about the signs of the *individual* eigenvalues, it doesn't turn itself off when there is no shear (all positive or negative eigenvalues). If $\nabla \cdot \tilde{\mathbf{u}} > 0$, the Boussinesq approximation can't distinguish between 2 positive eigenvalues and 1 negative eigenvalue or 2 negative eigenvalues and 1 positive eigenvalue. Finally, the above formula shows that for the Boussinesq approximation, we have

$$\Delta\lambda_i = \lambda_i - \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \ . \tag{34}$$

Thus there is an isotropic redistribution of the λ_i's and there is no enforcement that $|\Delta\lambda_i| \leq |\lambda_i|$. If we require a linear form for shear that can be written in tensor covariant form, then the Boussinesq approximation is the best we can do. Hopefully, the reader is convinced by now that there are better alternatives for modeling shear in the presence of dilatation..

4 Specialization to RZ geometries

Often we are interested in solving 2-dimensional problems in an RZ coordinate system where the velocity in the θ-direction is zero. The approach described above will work for this particular case, but one needs

to be careful, as small numerical errors in finding the λ_i and $\mathbf{e}_i$ could result in spurious (non-zero) values being generated in the $R\Theta$ or ΘZ components of S_{ij}^*. To get around this issue, we can take advantage of the decoupling of the θ-component to solve a 2x2 system instead of a 3x3 one and directly enforce orthogonality in the θ-direction. For example, in an RZ calculation, $\mathbf{S}$ has the form

$$
\begin{pmatrix}
S_{rr} & 0 & S_{rz} \\
0 & S_{\theta\theta} & 0 \\
S_{rz} & 0 & S_{zz}
\end{pmatrix} \ ,
\tag{35}
$$

where $S_{r\theta} = S_{\theta z} = 0$. $S_{\theta\theta}$ is proportinal to the velocity in the r-direction and is not in general zero.

Finding the eigenvalues results in the following characteristic equation

$$
(S_{\theta\theta} - \lambda) \left[\lambda^2 - \lambda (S_{rr} + S_{zz}) + \left(S_{rr} S_{zz} - S_{rz}^2 \right) \right] \ ,
\tag{36}
$$

from which we see directly that one of the eigenvalues, λ_θ, is $S_{\theta\theta}$. Solving the quadratic equation results in the other eigenvalues ($\lambda_\pm$) being

$$
\lambda_\pm = \frac{(S_{rr} + S_{zz}) \pm \sqrt{(S_{rr} - S_{zz})^2 + 4S_{rz}^2}}{2} \ .
\tag{37}
$$

It is obvious from the previous equation that the eigenvalues are all real, as they must be, since S_{ij} is symmetric.

For the eigenvalue given by λ_θ, the corresponding eigenvector must be $\mathbf{e}_\theta = (0, 1, 0)$. We can now solve the simpler 2x2 system given by

$$
\begin{pmatrix}
S_{rr} & S_{rz} \\
S_{rz} & S_{zz}
\end{pmatrix} \ ,
\tag{38}
$$

for the 2 other 2-dimensional orthogonal eigenvectors. We then convert these eigenvectors to 3-dimensional ones by inserting a 0 for the middle component. We now have our 3 eigenvalues and 3 orthogonal unit eigenvectors and we can use all of the above formulas and procedures discussed in the second section of the paper. By construction, we will not introduce any spurious values when we compute the necessary outerproducts to transform S_{ij}^* back into the original coordinate system.

5   Conclusions

We have shown that if one is going to model shear using a Boussinesq-like approximation in a RANS model, that one should NOT use the form given by

$$
D_{ij} = 2\mu_t \left(S_{ij} - \frac{1}{3} \nabla \cdot \tilde{\mathbf{u}} \ \delta_{ij} \right) \ .
\tag{39}
$$

Rather, one should use a form like

$$
D_{ij} = 2\mu_t S_{ij}^* \ ,
\tag{40}
$$

where the S_{ij}^* is constructed to be compression/expansion free and only represents true shear. The S_{ij}^* is determined by looking at the sign of $\nabla \cdot \tilde{\mathbf{u}}$ and the signs of the individual eigenvalues of $\mathbf{S}$, applying some reasonable constraints to the compression/expansion part of $\mathbf{S}$, subtracting to get the true shear part, and finally trnsforming from the principal basis back into the original coordinate system. This method will avoid the double counting between shear and dilatation that inadvertantly takes place if one uses the Boussinesq approximation. This method also significantly reduces realizability issues and eliminates the need for special switches like turning off the deviatoric stress in a shock. Any modifications/enhancements to the Boussinesq approximation should be made *after* switching to a form that uses S_{ij}^*, since we need to eliminate the double counting problem before we can tackle other deficiencies in the aprroximation.
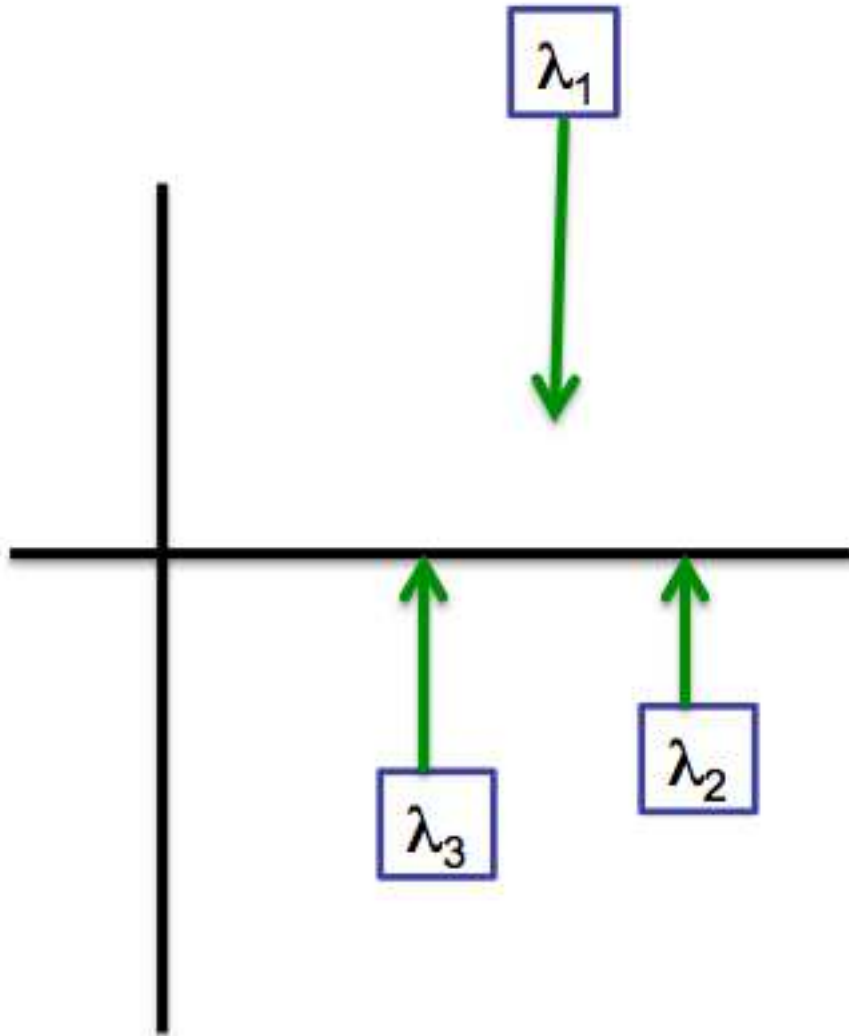
8

Figure 1: This is the case where $\nabla \cdot \tilde{\mathbf{u}} > 0$, λ_1 is positive, and λ_2 and λ_3 are negative. For this case we need to sort the eigenvalues so that $\lambda_1 \geq \lambda_2 \geq \lambda_3$. Here we adjust λ_2 and λ_3 upwards by their maximum allowed amounts, and then adjust λ_1 downward by the requisite amount to keep the sum of the changes identically zero.
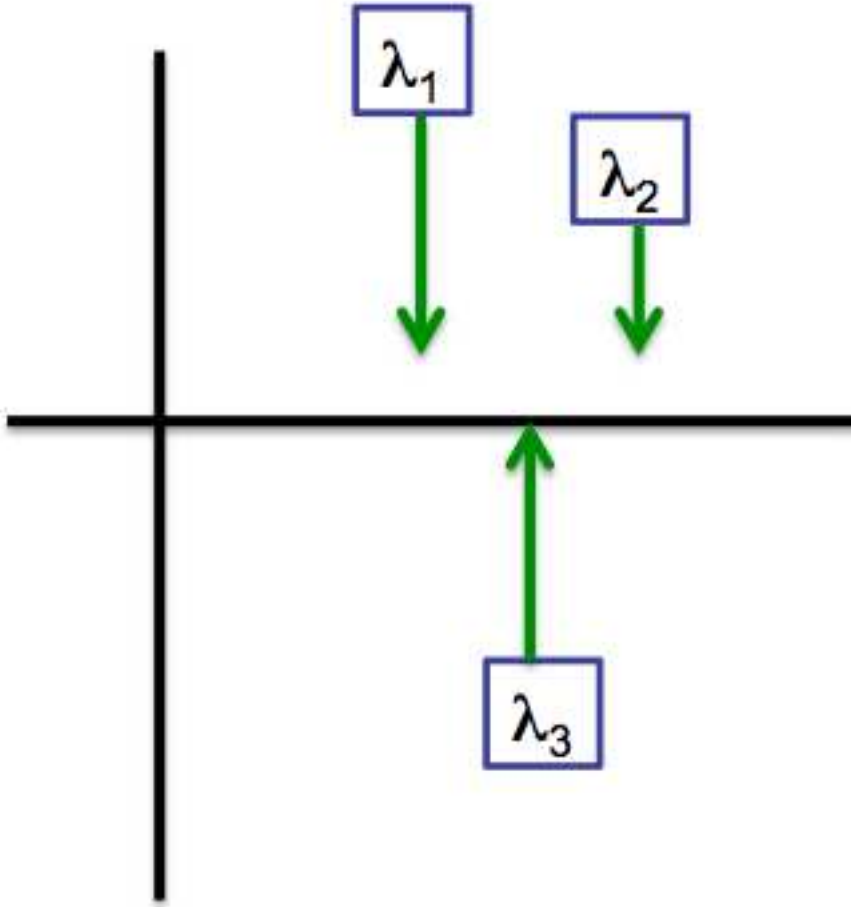
Figure 2: This is the case where $\nabla \cdot \tilde{\mathbf{u}} > 0$, λ_1 and λ_2 are positive and λ_3 is negative. For this case we need to sort the eigenvalues so that $\lambda_1 \geq \lambda_2 \geq \lambda_3$. Here we adjust λ_3 upwards by its maximum allowed amount, and then adjust λ_1 and λ_2 downward by their fractional amount of λ_3.
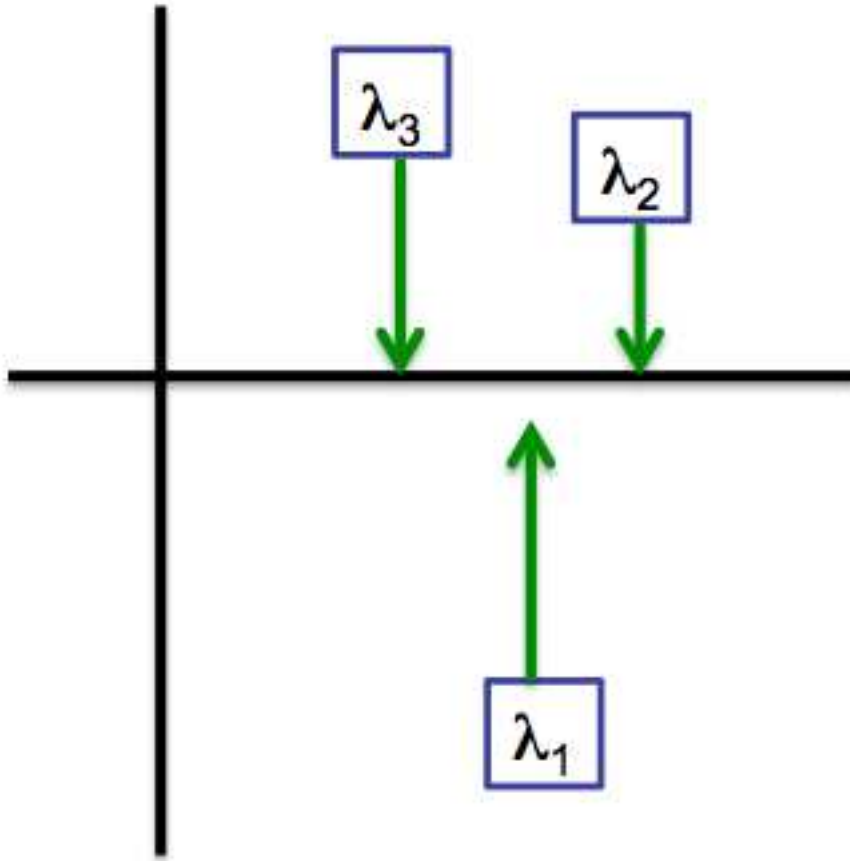
Figure 3: This is the case where $\nabla \cdot \tilde{\mathbf{u}} < 0$, λ_1 is negative and λ_2 and λ_3 are positive. For this case we need to sort the eigenvalues so that $\lambda_1 \leq \lambda_2 \leq \lambda_3$. Here we adjust λ_2 and λ_3 downwards by their maximum allowed amounts and then adjust λ_1 upward by the requisite amount to keep the sum of the changes identically zero.
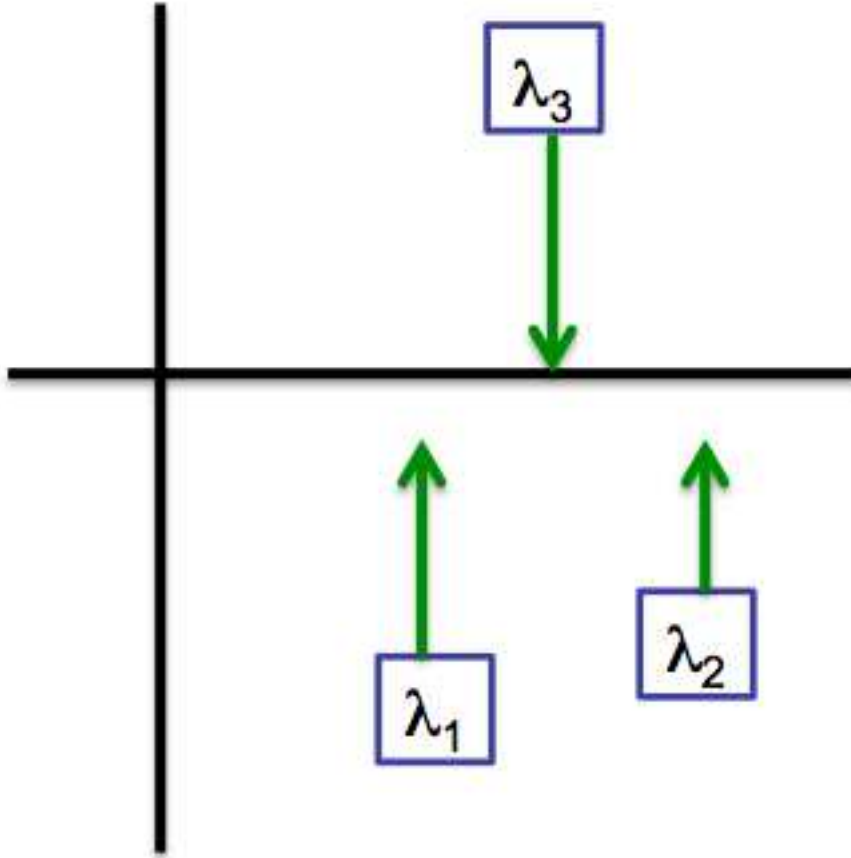
Figure 4: This is the case where $\nabla \cdot \tilde{\mathbf{u}} < 0$, λ_1 and λ_2 are negative, and $\lambda_3 > 0$. For this case we need to sort the eigenvalues so that $\lambda_1 \leq \lambda_2 \leq \lambda_3$. Here we adjust λ_3 downwards by its maximum allowed amount and then adjust λ_1 and λ_2 upward by their fractional amount of λ_3.